Does mutual knowledge of preferences lead to more Nash equilibrium play? Experimental evidence

Christoph Brunner[†] and T. Florian Kauffeldt[‡,⋆] and Hannes Rau[†]

Abstract

Nash equilibrium often does not seem to accurately predict behavior. In experimental game theory, it is usually assumed that the monetary payoffs in the game represent subjects' utilities. However, subjects may actually play a very different game. In this case, mutual knowledge of preferences may not be satisfied. In our experiment, we first elicit subjects' preferences over the monetary payoffs for all players. This allows us to identify equilibria in the games that subjects actually are playing. We then examine whether revealing other subjects' preferences leads to more equilibrium play and find that this information indeed has a significant effect. Furthermore, it turns out that subjects are more likely to play maxmin and maxmax strategies than Nash equilibrium strategies. This indicates that subjects strongly rely on heuristics when selecting a strategy.

Keywords: Behavioral Game Theory, Epistemic Game Theory, Nash Equilibrium, Incomplete Information Games, Strategic Ambiguity

JEL classifications: C91, C72

1. Introduction

People frequently show behavior that seems to be inconsistent with standard Nash equilibrium behavior such as cooperation in one-shot social dilemma situations. Nash equilibria in experiments are usually determined using the

[†] University of Heidelberg, Departement of Economics, Bergheimer Strasse 58, 69115 Heidelberg, Germany.

[‡] Heilbronn University, Departement of Economics, Max-Planck-Strasse 39, 74081 Heilbronn, Germany. [⋆] Corresponding author. Email: mail@florian-kauffeldt.com

monetary payoffs in the game. However, if subjects not only care about their *own* monetary payoffs, they may actually play a very different type of game. For example, instead of facing a social dilemma, the involved parties might actually play a coordination game (see Example 1 on page 5 for more details). In such situations, it is not clear whether there is common knowledge about preferences over the outcomes of the games. We study the impact of this aspect on equilibrium play.

Common (or, at least, mutual) knowledge about preferences is a core assumption in game theory. In the words of Polak (1999, p. 673):

> "In games of complete information, common knowledge of payoffs is usually taken to be implicit. Indeed, this is often taken to be the definition of complete information."[2]

Mutual or even common knowledge about preferences is not only assumed in traditional game theory, but also often in behavioral game theory. Most level-k models assume that payoffs are mutually known and that agents form beliefs about other agents' play based on this information (see, e.g., Costa-Gomes et al., 2001). Other approaches such as Quantal Response Equilibrium[3] introduced by McKelvey and Palfrey (1995) incorporate a stochastic element that can be interpreted as uncertainty about other players' preferences. Our findings might also be relevant for these concepts.

Despite the ubiquity of the (implicit) assumption of mutually or commonly known preferences, there is little empirical evidence about the degree to which it affects the reliability of the Nash prediction. However, previous experimental research suggests that it should not be taken for granted. For example, Healy (2011) finds that subjects fail to accurately predict other subjects' preferences over possible outcomes in normal-form 2×2 games. The purpose of the experiment reported in this paper is to test whether mutual knowledge of preferences is important for the Nash prediction.[4]

[2]This means that "complete information" cannot be part of the rules of the game (*the game-form*) because it involves assumptions about knowledge of individual preferences.

[3]In Quantal Response Equilibrium, there is an error term in players' payoff functions whose distribution is assumed to be known.

[4]Notice that we examine one-shot dominance-solvable 2 x 2 games. In the context of repeated games, it has been studied whether subjects are able to learn the preferences of their opponents (see, e.g., Oechssler and Schipper, 2003). However, it is not the purpose of our experiment to examine learning in games. Repeated games require repeated interaction and are much more complex to analyze than one-shot games.

Our results can be summarized as follows: (1) subjects are indeed significantly more likely to play a Nash equilibrium strategy when they are informed about their opponents' preferences over the possible outcomes of the game. When preferences are not mutually known, the frequency of equilibrium play is rather low. (2) A strategy is more likely to be played when it cannot lead to the lowest payoffs (*maxmin strategy*) or when it can lead to the highest one (*maxmax strategy*). Furthermore, maxmin and maxmax strategies predict behavior better than Nash equilibrium strategies, especially when preferences are not mutually known.

Result (1) shows that subjects not only fail to accurately predict other players' preferences, the lack of such information also significantly affects their behavior. Whenever it is unlikely that players know each other's preferences and some players have no strictly dominant strategy, it might therefore be advisable to use a more general equilibrium concept. Following Polak (1999), we may view a situation where preferences are not mutually known as a game with incomplete information. Such a situation can then be modeled as a Bayesian game (Harsanyi, 1967-1968).[5]

Result (2) suggests that subjects largely rely on heuristics rather than on strategic considerations. The reason may be that subjects face two types of uncertainty in the experiment: there is uncertainty about the other player's payoff function in the baseline treatment. In the information treatment, subjects are informed about the other player's preferences but they still may not believe that he is rational. Uncertainty about opponent's rationality and/or payoff function can lead to uncertainty about the other agents' strategy choices with unknown probabilities. Uncertainty about probabilities (*ambiguity*) can affect peoples' behavior, as Ellsberg (1961) illustrated.[6] The strategic ambiguity model of Eichberger and Kelsey (2014) shows that maxmin and/or maxmax strategies can be a best response to strategic am-

[5]Players with different preferences can be thought of as different types and it is then assumed that the prior distribution of types is commonly known. This approach has been taken in various fields. In auction theory, e.g., the assumption that all bidders are risk neutral and that this is commonly known has been relaxed (see, e.g., Hu and Zou, 2015).

[6]When people face ambiguity, they frequently do not behave as if they were governed by subjective probabilities.

biguity.[7]

Theoretically, in the tested dominance-solvable 2×2 games, mutual knowledge of payoff functions along with mutual knowledge of rationality suffices to ensure that agents will play a Nash equilibrium.[8] To see this, suppose one player (called "D") has a strictly dominant strategy. Given that D is assumed to know his own payoff function and is rational, D will play his dominant strategy. The other player (called "ND") believes that D is rational and that he has a strictly dominant strategy. Therefore, ND believes that D will play this strategy. Since ND is himself assumed to be rational and to know his own payoff function, ND will play a best response to D's dominant strategy.

1.1. The experiment

In stage 1 of the experiment, we elicit subjects' preferences over monetary payoff pairs (they will be referred to as "payment pairs"). The same payment pairs are then used to construct eight different 2×2 games (or more precisely eight different game-forms). In stage 2, each subject plays four out of the eight games exactly once. We ran two waves of experiments so that all games are played. Our design allows us to avoid the assumption that subjects only care about their *own* monetary payments. Instead, we can use the preferences elicited in stage 1 to describe the game that our subjects play.[9]

[7]This model allows for optimistic responses to strategic ambiguity. Most other strategic ambiguity models such as those of Lo (1996), Eichberger and Kelsey (2000), and Lehrer (2012) assume ambiguity-averse behavior. While these models can explain maxmin behavior, they cannot rationalize maxmax behavior.

[8]See Aumann and Brandenburger (1995) for sufficient conditions that ensure Nash equilibrium in general normal-form 2×2 games. Notice that there is a difference between "knowledge" and "(probability one) belief". Roughly, "knowledge" refers to true belief justified by either direct observation or logical deduction, whereas "belief" may be false. Therefore, it would be more accurate to assume that players believe that others are rational with probability one.

[9]We maintain the assumption that preferences depend only on players' monetary payments. That is, the specific game-form, other subjects' preferences, or any other factors have no effect on subjects' ordinal ranking of payment pairs. Of course, this is to some degree a consequentialist approach and consequentialism has been criticized in the literature repeatedly. In Section 3.4, we will discuss evidence suggesting that such considerations do not play an important role in the games used in this study. However, we cannot completely exclude that violations of consequentialism might have caused some noise and led to a systematic downward bias so that the "true" treatment effect is even higher.

This will be illustrated with the help of Example 1 below, which corresponds to one of the games played in the experiment.

Example 1. *Consider the prisoner's-dilemma-type game-form in Figure 1. The numbers in the matrix correspond to the amount of money paid to the players, where the first number is the row player's payment and the second number is the column player's payment.*

	L	R
U	$4, 4$	$8, 3$
D	$3, 8$	$7, 7$

Figure 1: Prisoner's-dilemma-type game-form

Depending on players' preferences over the payment pairs various games can be induced by the game-form in Example 1. Let r be the row and c be the column player and denote the payment pairs by $(x_r, x_c) \in \mathbb{R}^2$. Suppose player i's ($i \in \{r, c\}$) preferences over payment pairs are represented by a function $v_i : \mathbb{R}^2 \to \mathbb{R}$. In general, the games induced by the game-form in Example 1 take the following form:

	L	R
U	$v_r(4, 4), v_c(4, 4)$	$v_r(8, 3), v_c(8, 3)$
D	$v_r(3, 8), v_c(3, 8)$	$v_r(7, 7), v_c(7, 7)$

Figure 2: Induced games by the game-form in Example 1

The induced game is only a prisoner's dilemma game if the players mainly care about their own payoffs, e.g., if $v_i(x_r, x_c) = x_i$ for $i \in \{r, c\}$. In this case, the game has only one Nash equilibrium (U, L), i.e., everyone defects. On the other hand, the game that results if players have other-regarding preferences can be very different. For instance, suppose players' preferences are represented by the following utility function $v_i(x_r, x_c) = \min\{x_r, x_c\}$ for $i \in \{r, c\}$. Then, (U, L) as well as (D, R), i.e. mutual cooperation, are Nash equilibria.

In this paper, whenever we refer to a "Nash equilibrium", we refer to the Nash equilibria of the induced game using the preferences elicited in stage 1 of the experiment. We focus on those situations in which a unique pure Nash equilibrium exists (according to the reported preferences): one player has a strictly dominant strategy and the other player has a non-dominant unique pure Nash equilibrium strategy in the induced game.[10] Consequently, we consider situations in which subjects' opponents have a strictly dominant strategy. In the baseline treatment, the reported preferences are not revealed. Hence, subjects cannot be certain that their opponents have a dominant strategy.

For example, suppose the row player in the induced game above is selfish. His pure strategy U is then strictly dominant. A column player who prefers $(4, 4)$ to $(8, 3)$ and $(7, 7)$ to $(3, 8)$ then has a unique equilibrium strategy that is not dominant: L. In treatment baseline, such a column player may not be sure whether row is selfish or not and might therefore occasionally play R rather than L.

In our second treatment (called "info"), the column player can see that row has a strictly dominant strategy and might therefore play the unique equilibrium strategy L more often. Intuitively, this logic can explain our first result that subjects are more likely to play a Nash equilibrium strategy in treatment info compared to treatment baseline. Furthermore, if a subject is uncertain about the strategy choice of his opponent, then, depending on his attitude towards uncertainty, he will try to avoid the lowest ranked payment pair (maxmin), or, to reach the highest ranked one (maxmax). This intuition explains our second result.

1.2. Related literature

The papers closest to ours are Healy (2011) and recent working papers by Wolff (2014) and Attanasi et al. (2016).

[10] In our experiment, we only ask subjects to rank payment pairs ordinally. Eliciting a cardinal ranking of payment pairs would require a more complicated procedure that some subjects might fail to understand. It is not obvious that subjects can reliably assign a cardinal utility to each payment pair. As a result, we cannot compute Nash equilibria in mixed strategies for the induced games. Moreover, we will exclude the decisions of subjects who have a strictly or weakly dominant strategy in the induced game. Information about their opponent's preferences is not necessary for those subjects to compute a best response and as a result, information about the other player's preferences should not be expected to have an effect on behavior.

Healy examines whether the sufficient conditions for Nash equilibrium identified by Aumann and Brandenburger (1995) are satisfied when subjects play normal-form 2×2 games in the laboratory. For that purpose, subjects first choose a strategy and then report their beliefs about behavior and preferences of their opponent. Subjects' own preferences and rationality are also measured. Healy finds that there are only very few instances where all conditions are satisfied. Focusing on mutual knowledge of preferences, he finds that both players correctly predict how their opponent ordinally ranked the payment pairs in only 64% of games played. Healy concludes that "The failure of Nash equilibrium stems in a large part from the failure of subjects to agree on the game they are playing."

Since mutual knowledge of preferences is one of three conditions that are together sufficient for Nash equilibrium in 2×2 games (see Aumann and Brandenburger, 1995) and since the other two are also not fully satisfied in Healy's experiment, it is difficult to assess the impact of the failure of mutual knowledge of preferences on equilibrium play in isolation. By introducing a treatment in which information about the opponent's preferences is directly revealed, we can identify the impact of mutual knowledge on equilibrium play by holding all other factors constant.

Wolff (2014) studies behavior in three-person sequential public good games. In contrast to our experiment, he does not reveal subjects' preferences over the material outcomes. Instead, he elicits subjects' best-response correspondences to the contributions of the other players. In one of his treatments, these are then revealed to all group members. This information has a much smaller effect on the frequency of equilibrium play compared to the treatment effect in our experiment.

Revealing best-response correspondences is obviously not sufficient for subjects to be able to predict how much their opponents will contribute: Wolff measures beliefs about others' contributions to the public good and finds that subjects tend to overestimate these. As a result, they often fail to play an equilibrium strategy even though their contributions tend to be consistent with their beliefs and their own reported best-responses. As opposed to the dominance-solvable 2×2 games that we study, several iterations of alternating best responses are required in Wolff's experiment to compute the Nash equilibrium. Some subjects might not be able to do so.

Attanasi et al. (2016) also argue that when subjects have belief-dependent or other-regarding preferences, they are actually playing a game of incomplete information. In their experiment, subjects form beliefs about their opponent's type (e.g., selfish or prosocial) and choose their strategy based on these belief. Attanasi et al. then test whether revealing information about opponent's preferences and beliefs changes behavior in a Mini Trust Game. They find that first movers are more likely to transfer the money when they face a non-selfish trustee ("guilt-averse" trustee) and vice versa. The Mini Trust Game can be considered as a 2×2 coordination game with two pure equilibria (trust, share) and (not trust, not share). Subjects clearly coordinate better on one of these two equilibria when belief-dependent preferences are disclosed. While this result points in a similar direction as our results, Attanasi et al. do not systematically test the impact of mutual knowledge of preferences on the Nash prediction. In particular, the second movers can observe the decisions of the first movers. Therefore, the preferences of the first movers are not relevant for their strategy choices.

This paper is organized as follows. The next section describes the experimental design. We then present our results, and conclude in Section 4. The appendix provides additional information about the experiment.

2. Experimental design

Our experiment consists of two treatments (called "baseline" and "info") with two stages each. In the first stage of both treatments, we elicit subjects' preferences over eight different payment pairs. These payment pairs are then used to construct eight different 2×2 games. In stage 2, each subject plays four out of the eight games exactly once. We ran two waves of experiments. Subjects played Game 1 to 4 in wave 1 and Game 5 to 8 in wave 2.[11] In treatment "info", subjects can see their opponent's ordinal ranking of the four payment pairs used in the current game, whereas in treatment "baseline", this information is not disclosed.

2.1. Stage 1 of the experiment

Stage 1 is identical in both treatments. Subjects are asked to create an ordinal ranking over the following set X_{row} of eight payment pairs (x_r, x_c):

$$X_{row} = \{(8,3),\ (7,7),\ (5,8),\ (4,4),\ (6,2),\ (3,8),\ (3,3),\ (2,2)\} \qquad (1)$$

The first number, x_r, corresponds to the amount of money (in Euros) paid to the decision-maker in the role of a row player. The second number, x_c, is

[11] The games are described in detail in section 2.2.

paid to some other subject in the role of a column player (the "recipient").[12] Subjects are informed that they will not interact with the recipient in any other way in either stage of the experiment.

The order in which the payment pairs appear on the screen was randomly determined beforehand and remains constant in all sessions. Subjects rank the payment pairs by assigning a number between one and eight to each pair, where lower numbers indicate a higher preference. The same number can be assigned to multiple payment pairs, thus allowing for indifference.

In treatment info, subjects are told that their rankings would be disclosed to other participants at a later stage of the experiment.[13] In treatment baseline, we made it clear that this information would not be revealed. We will explain at the end of this section how the elicitation of preferences was incentivized. After subjects confirm their ranking, they proceed to stage 2, in which they play four one-shot 2×2 games.

2.2. Stage 2 of the experiment

We ran two waves of experiments. In the first wave, subjects played the games in Figure 3 (all numbers are payments in Euro). In the second wave, they played the games in Figure 4. All games were constructed using the same eight payment pairs, see set X_{row} defined in (1).

We made sure that the games exhibit some diversity with respect to the number of pure strategy Nash equilibria under the assumption that subjects are selfish payment maximizers. The eight games were selected on the basis of two key criteria that seem to play an important role in the context of our study:

(i) # players, who have a strictly dominant strategy (0, 1 or 2) and

(ii) # pure Nash equilibria (0, 1 or 2).

Both criteria were determined for the case where preferences correspond to monetary payoffs (i.e., on the basis of the game-forms). 2×2 games can

[12]Subjects who were assigned the role of a column player ranked the same payment pairs but the first number corresponds to the other player's payoff. Rewriting X_{row} for column players such that the first number corresponds to the column player's payment and the second to the row player's, we obtain $X_{column} = \{(8,3), (7,7), (8,5), (4,4), (2,6), (3,8), (3,3), (2,2)\}$.

[13]We will discuss the possibility that subjects might strategically misrepresent their preferences in the results section.

be grouped into 6 categories based on these two criteria (some combinations are not possible, e.g., 2 players with strictly dominant strategies and 2 Nash equilibria). Games with more than 2 pure equilibria are unlikely to offer valuable insights for our analysis because they are not expected to generate many relevant observations. We first run wave 1 of the experiment, then we selected the games of wave 2 so that we have at least one game of each of the 6 categories. Furthermore, we wanted to cover most of the 2×2 games that are frequently used in experimental economics (e.g., Prisoners' Dilemma, Matching Pennies, and Battle of Sexes).

Game 1

	L	R
U	$4,4$	$8,3$
D	$3,8$	$7,7$

Game 3

	L	R
U	$4,4$	$8,3$
D	$3,3$	$7,7$

Game 2

	L	R
U	$5,8$	$7,7$
D	$6,2$	$3,3$

Game 4

	L	R
U	$8,3$	$2,2$
D	$7,7$	$3,8$

Figure 3: Games in wave 1

Game 5

	L	R
U	$3,8$	$8,3$
D	$3,3$	$7,7$

Game 7

	L	R
U	$8,3$	$6,2$
D	$7,7$	$5,8$

Game 6

	L	R
U	$8,3$	$2,2$
D	$2,2$	$3,8$

Game 8

	L	R
U	$3,3$	$8,3$
D	$2,2$	$7,7$

Figure 4: Games in wave 2

2.3. Experimental procedure, treatment information and incentives

In both treatments, subjects can see how they ranked the four payment pairs of the currently played game. This information is displayed by assigning 1-4 stars to each outcome, where more stars indicate a better outcome. In treatment info, subjects are shown both their own *and* their opponent's ranking in matrix-form (see Figure 5). Just like in the payment matrix, the first entry corresponds to the subject's own ranking while the second entry reveals the opponent's ranking. In treatment baseline, subjects are shown the same rankings matrix but this matrix only contains their own rankings.

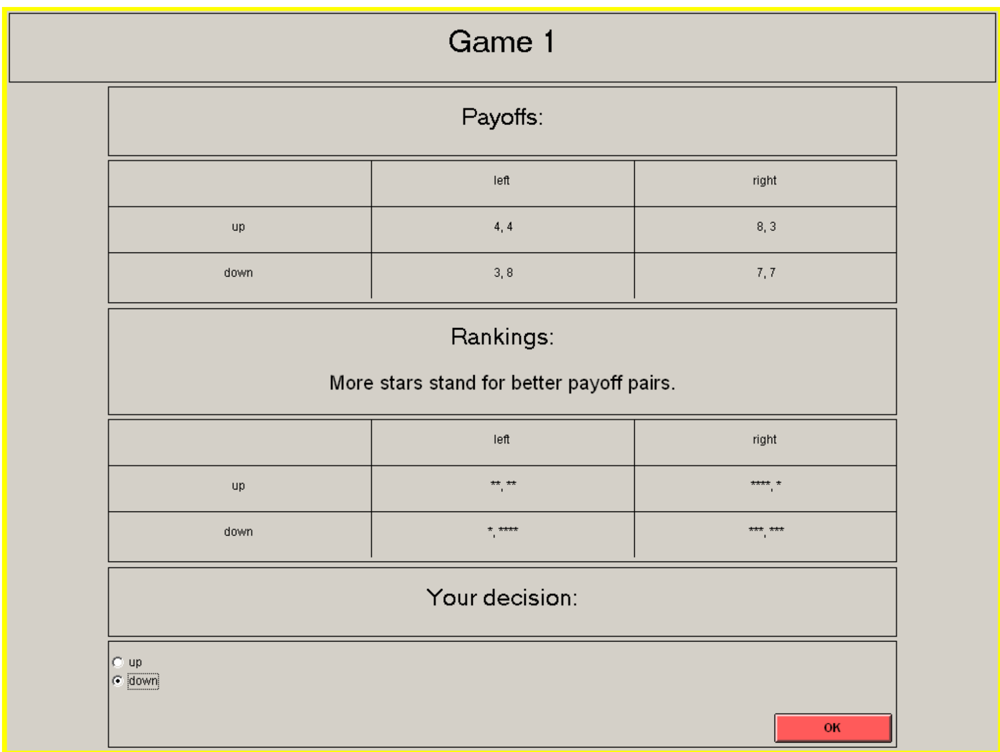

Figure 5: Information screen

All subjects play each of the four games of their wave exactly once, each time against a different anonymous opponent. Games are played one after another and feedback about the outcome is only provided at the end of the experiment when subjects are paid, but not while subjects still make decisions.

In both treatments, each subject is paid for exactly one of his decisions, which is randomly selected at the end of the experiment. If a decision from stage 1 is chosen, two of the eight payment pairs from the set X_{row} are randomly selected. The row subject is then paid the first number, x_r, of the payment pair that he ranked more highly in stage 1. The second number, x_c, is paid to some other column subject. In order to avoid reciprocity considerations, we made it clear that the second number is paid to a subject with whom subjects will not interact in the second stage of the experiment. Column subjects are paid in a similar manner.

The probability that stage 1 is paid is $\frac{7}{8}$ while stage 2 is paid with a probability of $\frac{1}{8}$. These probabilities are consistent with selecting each of the $\binom{8}{2}$ possible pairs of payment pairs and each of the four decisions made in stage 2 with equal probability. Paying stage 1 with a substantially higher probability also reduces the odds that subjects might misrepresent their preferences. This issue will be discussed in more detail in Section 3.4.

Subjects were given printed instructions and they could only participate after successfully answering several test questions. Test questions as well as the rest of the experiment were programmed using Z-Tree (Fischbacher, 2007). All sessions of the experiment were conducted at the AWI-Lab of the University of Heidelberg. Subjects from all fields of study were recruited using Orsee (Greiner, 2015). Fewer than half of the subjects were economics students. Sessions lasted about 40-50 minutes on average. The following table summarizes the number of participants per session as well as average payments:

Table 1: Summary of treatment information

Treatment	Wave	Sessions	Subjects	Average payment
baseline	1	9	97	€ 12.02
baseline	2	7	91	€ 10.54
info	1	8	95	€ 11.78
info	2	7	85	€ 11.41

Decisions made by subjects who made more than 10 mistakes when answering test questions are excluded from the data (including Table 1).[14]

3. Results

In this section, we first characterize subjects' preferences as measured in stage 1 of the experiment. We then present the main treatment effect: subjects are significantly more likely to play their unique equilibrium strategy in treatment info than in treatment baseline. This effect can be observed in 6 of the 8 games. Subsequently, we show that maxmin and maxmax strategies are more likely to be played in both treatments. We argue that it is unlikely that subjects misrepresent their true preferences or that many preferences changed when subjects are shown their opponents' preferences.

3.1. Characterization of measured preferences

In stage 1 of the experiment, we elicit subjects' preferences over the payment pairs $(x_r, x_c) \in X_{row}$ defined in equation (1). Tables 2 and 3 show the ordinal rankings reported by at least two subjects who were assigned the role of a row and column player respectively. Payment pairs that are assigned a lower number are preferred to payment pairs with a higher number.

[14]The main treatment effect (Table 7) is still significant when these 10 subjects are included. In treatment baseline, 2 subjects made more than 10 mistakes, in treatment info, there were 8 such subjects. It is not plausible that the decisions of the excluded subjects affected other subjects' decisions since all of our games are simultaneous games and subjects were not informed about the decisions of their opponents during the experiment.

Table 2: Preferences reported by at least two subjects who were assigned the role of a row player, both treatments. Smaller numbers are assigned to better ranked payment pairs.

(8,3)	(7,7)	(5,8)	(4,4)	(6,2)	(3,8)	(3,3)	(2,2)	n
1	2	4	5	3	6	7	8	63
1	2	4	5	3	7	6	8	15
2	1	4	5	3	6	7	8	15
1	2	4	5	3	6	6	8	12
2	1	3	5	4	6	7	8	10
2	1	3	6	4	5	7	8	7
1	2	3	5	4	6	7	8	5
2	1	3	4	5	6	7	8	5
3	1	2	5	6	4	7	8	3
3	1	2	5	5	3	7	8	3
1	2	3	5	3	6	7	8	2
1	2	3	4	5	6	7	8	2
3	1	2	6	5	4	7	8	2
3	1	2	4	5	6	7	8	2
1	1	4	5	3	6	7	8	2
1	2	5	4	3	7	6	8	2
1	2	4	6	3	5	7	8	2

Table 3: Preferences reported by at least two subjects who were assigned the role of a column player, both treatments. Smaller numbers are assigned to better ranked payment pairs.

(8,3)	(7,7)	(8,5)	(4,4)	(2,6)	(3,8)	(3,3)	(2,2)	n
2	3	1	4	7	5	6	8	68
3	2	1	4	7	5	6	8	14
3	1	2	4	7	5	6	8	14
3	1	2	5	6	4	7	8	8
1	3	1	4	7	5	5	7	5
3	1	2	4	8	6	5	7	5
1	3	1	4	8	6	5	7	4
3	2	1	5	7	4	6	8	4
2	3	1	4	7	5	5	7	4
3	2	1	5	6	4	7	8	3
1	3	1	4	7	5	6	8	3
4	1	2	3	6	5	7	8	2
1	3	2	4	8	6	5	7	2
3	1	2	4	6	5	7	8	2
2	3	1	4	6	5	7	8	2
3	2	1	4	8	6	5	7	2
3	1	2	4	8	7	5	6	2

To characterize subjects' preferences, we introduce four properties: pareto-efficiency, strict pareto efficiency, maximization of own payoff, and maximization of total payoff. These properties are defined as follows:

Definition 1 (Pareto efficiency)**.** A subject's preferences $\succsim$ on X are said to satisfy *pareto-efficiency* if, for all $x, y \in X_{row}$, $x \succ y$ whenever $x_r \geq y_r$ and $x_c \geq y_c$ with at least one inequality strict.

Definition 2 (Strict pareto efficiency)**.** A subject's preferences $\succsim$ on X are said to satisfy *strict pareto-efficiency* if, for all $x, y \in X_{row}$, $x \succ y$ whenever $x_r > y_r$ and $x_c > y_c$.

Definition 3 (Own payoff maximization)**.** A row (column) subject is said to *maximize his own payoff* if, for all $x, y \in X_{row}$ ($x, y \in X_{column}$), $x \succ y$ whenever $x_r > y_r$ ($x_c > y_c$).

Definition 4 (Total payoff maximization)**.** A subject is said to *maximize total payoff* if, for all $x, y \in X_{row}$, $x \succ y$ whenever $x_r + x_c > y_r + y_c$.

Table 4 shows the fraction of subjects whose preferences are consistent with the properties defined above.

Table 4: Measured preferences

Treatment	Pareto efficiency	Strict pareto efficiency	Own payoff max.	Total payoff max.	n
Pooled	70.9%	90.2%	48.6%	4.6%	368
Baseline	71.8%	90.4%	46.8%	4.3%	188
Info	70.0%	90.0%	50.6%	5.0%	180

Preferences that satisfy pareto efficiency or own payoff maximization must also satisfy strict pareto efficiency. The vast majority of subjects report preferences that are consistent with strict pareto efficiency. Table 5 further classifies those preferences. Clearly, most preferences that satisfy strict pareto efficiency also satisfy either pareto efficiency or own payoff maximization or both. Only 6.9% of the preferences that satisfy strict pareto efficiency are not consistent with either pareto efficiency or own payoff maximization (pooled data). Notice also that preferences that satisfy total payoff maximization must simultaneously satisfy pareto efficiency.

Table 5: Preferences that satisfy strict pareto efficiency

Treatment	Pareto	Own payoff max.	Pareto and Own payoff max	Only strict pareto	n
Pooled	78.6%	53.9%	39.5%	6.9%	332
Baseline	79.4%	51.8%	37.6%	6.5%	170
Info	77.8%	56.2%	41.4%	7.4%	162

3.2. Nash equilibrium play

Our first hypothesis is that subject behavior is more consistent with the Nash equilibrium when preferences are mutually known. We test this hypothesis by using two different subsets of our data, which are depicted in Figure 6. Notice that we use the preferences elicited in stage 1 to identify dominant and equilibrium strategies.

There are a total of 368 subjects who participated in the experiment. Since each subject played four games (=four decisions) in stage 2, we have data on 1472 individual decisions, 752 in treatment baseline and 720 in treatment info. Only one type of strategic situation is of interest: a unique pure Nash equilibrium where one player has a unique equilibrium strategy that is neither weakly nor strictly dominant.[15] Therefore, we exclude those decisions where both strategies are played with strictly positive probability in some Nash equilibrium. Furthermore, we exclude those decisions where the equilibrium strategy is weakly or strictly dominant. In such a situation, the best response does not depend on the other player's strategy and therefore,

[15] Notice that this is only possible when the other player has a strictly dominant strategy.

it should not matter whether or not the other players' preferences are known. This leaves us with 279 relevant individual decisions, 140 in treatment baseline and 139 in treatment info. In all of the 279 games, the subject whose decision we study has a unique pure equilibrium strategy and that subject's opponent has a strictly dominant strategy. We test our main hypothesis using these 279 observations and will refer to the corresponding subset of our data as **"all subjects"**.

We also consider a subset of the set "all subjects" which no longer includes the decisions made by subjects who played a strictly dominated strategy in at least one of the four games. Either the preferences that these subjects reported in stage 1 do not reflect their true preferences or they are not rational in the sense that their choice in stage 2 is inconsistent with their reported preferences. Table 6 shows that approximately one fourth of our subjects violate strict dominance at least once. Removing the choices made by inconsistent subjects therefore further reduces the number of observations to 226 individual decisions, 115 in treatment baseline and 111 in treatment info. We will refer to this subset of our data as **"consistent subjects only".**

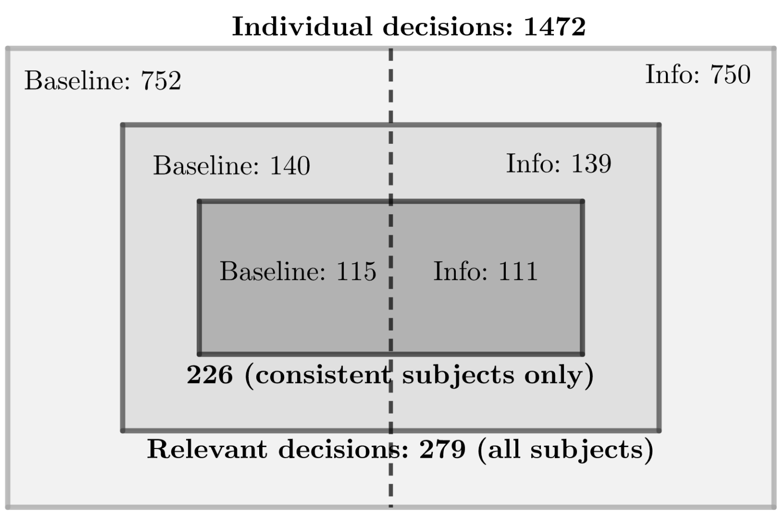

Figure 6: Relevant observations for equilibrium analysis

Table 6: Violations of strict dominance

Treatment	Subjects	Games played	Games with dominant strategy	Dominated strategy played	Subjects who played dominated strategy at least once
Baseline	188	752	280	23.2%	26.1%
Info	180	720	295	24.4%	29.4%

Figure 7 shows that subjects play an equilibrium strategy more often in treatment info than in treatment baseline.

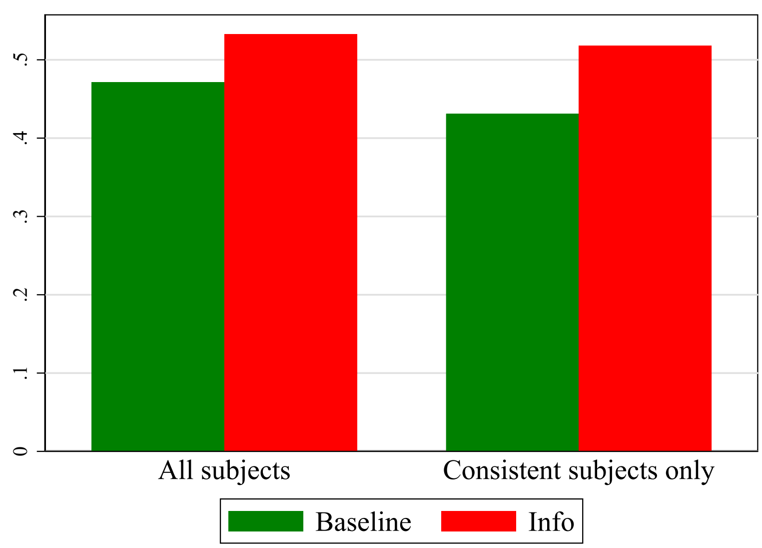

Figure 7: Frequencies of played unique equilibrium strategies

To test whether these differences are significant, we run a logit regression. The dependent variable "equilibrium strategy played" assumes a value of 1 if a subject plays the unique equilibrium strategy and 0 otherwise. We include an intercept as well as a dummy variable, which assumes a value of 1 if the observation is generated in treatment info and 0 otherwise. These results are shown in Table 7. The treatment effect is significant indicating that informing subjects about their opponents' preferences leads to a higher frequency of equilibrium play. Furthermore, the treatment effect is comparable when we only use the decisions made by these consistent subjects, even though the number of observations is reduced by approximately 20%.

Table 7: Logit regression "equilibrium strategy played", robust standard errors clustered by subject

Dependent variable: equilibrium strategy played	All Subjects	Consistent subjects only
info	0.54**	0.60**
	(0.26)	(0.29)
constant	-0.41**	-0.41**
	(0.18)	(0.20)
n	279	226
Clusters	212	166
Pseudo R^2	0.013	0.016

** significant at 5% level

Notes: Marginal fixed effects of change of variable "info" from 0 to 1 are 0.13** for the specification "All subjects" and 0.15** for "Consistent subjects only"

We also test whether there is a significant treatment effect using a two-tailed two-sample Wilcoxon rank-sum test. The dependent variable is the frequency with which a subject played an equilibrium strategy. Each subject who plays at least one game where the subject has a unique equilibrium strategy that is not weakly or strictly dominant counts as one observation. We run the same test for all subjects and for consistent subjects only. When using all (only consistent) subjects, we have 107 (87) observations in treatment baseline and 105 (79) in treatment info. The null hypothesis that the distribution of the frequency of equilibrium play is the same in both treatments can be rejected regardless of which data set we use.[16]

Result 1. *Subjects are more likely to play their unique Nash equilibrium strategy when preferences are mutually known.*

[16] p=0.083 using all subjects, p=0.086 using consistent subjects only.

As a robustness check, we also compute the frequency of equilibrium play for each game separately. These results are shown in Figure 8 for all subjects and in Figure 9 for consistent subjects only. Information about the number of relevant decisions for each game can be found in the tables A.10 and A.12 in the appendix. Regardless of which subset of our data we use, the frequency of equilibrium play is higher in treatment info than in treatment baseline for every game[17] except for games 5 and 6.

At first glance, subject behavior in Game 5 appears to be surprising: there is less equilibrium play in treatment info than in treatment baseline. A detailed check shows that all subjects, who did not take the equilibrium stategy in treatment info, were column players who played strategy R. The row players had the strictly dominant strategy U. Our second main result in Section 3.3 shows that many subjects followed a heuristic approach by selecting maxmax and/or maxmin strategies. Game 5 exhibits a special feature: the equilibrium and maxmax/maxmin strategy especially often fall apart. In several cases the (non-equlibrium) strategy R is both the maxmax and the maxmin strategy. These cases occur considerably more often in treatment info than in treatment baseline: in treatment info, 6 out of 10 subjects who violated the equilibrium prediction faced such a situation, while this is only the case for 1 out of 3 subjects in treatment baseline.

In Game 6 the frequency of equilibrium play is zero in both treatments as Figure 8 shows. This is in line with what we have expected: recall that Game 6 is a "Battle of Sexes"-type game-form and we expected that the game that subjects actually play (the induced game) is in most cases a "Battle of Sexes"-type game. Consequently, the situation that one subject has a unique non-dominant equilibrium strategy occurs very rarely here (for consistent subjects, we only have 5 relevant cases in both treatments together). Nonetheless, it makes sense to incorporate this situation in our analysis. First, we intend to provide a comprehensive analysis of the most popular games used in the experimental literature as described in the introduction. Second, we wanted to test whether the outcome meets our expectations.

[17]Using a Fisher exact test, this difference is significant at the 5% level for Game 3, when we use all subjects. We have more observations for Game 3 than for any other game. In Game 3, it occurred particularly often that one subject had a strictly dominant equilibrium strategy while the other subject did not have a strictly or weakly dominant strategy. Details of these tests can be found in the appendix (tables A.10 and A.12).

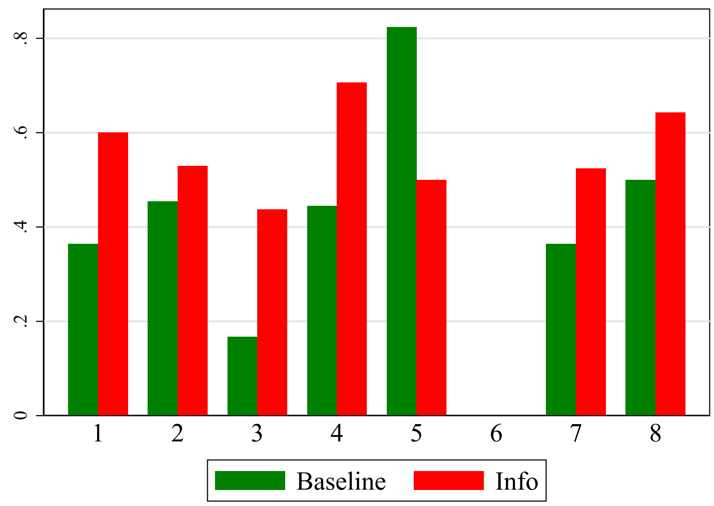

Figure 8: Frequency of equilibrium play by game, all subjects

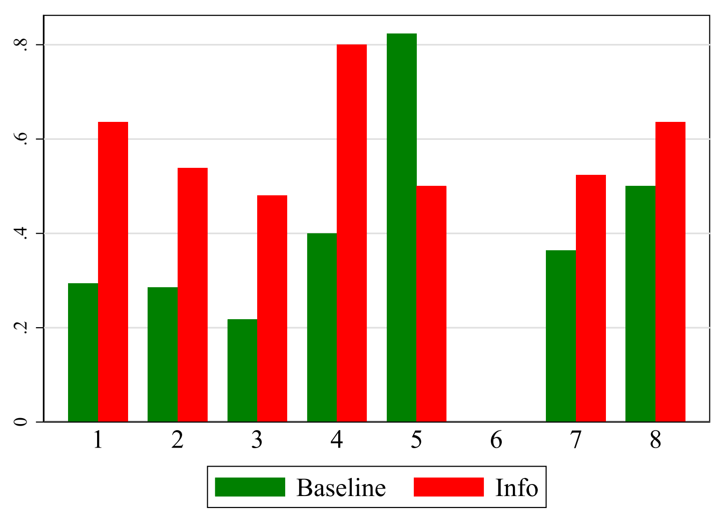

Figure 9: Frequency of equilibrium play by game, consistent subjects only

3.3. Maxmin and maxmax strategy play

As outlined in the introduction, playing a maxmin or a maxmax strategy can be a response to strategic uncertainty. Therefore, our second hypothesis is that a strategy is more likely to be played when it is a maxmin and/or a maxmax strategy. Subjects may use these strategies when they are uncertain about other players' payoff functions and/or other players' rationality. In treatment baseline, subjects face both types of uncertainty, whereas the uncertainty about other players' payoffs is removed in treatment info. Since there is some uncertainty in both treatments, we would expect a strategy to be played more often if it is a maxmin or a maxmax strategy in both treatments. Both effects are expected to be stronger in treatment baseline compared to treatment info.

We test these conjectures by running a conditional logit regression. An observation corresponds to a pure strategy. The dependent variable ("played") assumes a value of 1 if a strategy is played and 0 otherwise. Three independent variables are used to characterize each strategy: "equilibrium" indicates whether a strategy is a Nash equilibrium strategy. "maxmax" assumes a value of 1 if a strategy can lead to a most highly ranked payment pair. "maxmin" indicates whether a strategy can result in the realization of a lowest ranked payment pair (maxmin = 0 if that is the case, maxmin = 1 otherwise).

We only use decisions made by subjects who never played a strictly dominated strategy. Table 8 shows that whether or not a strategy is a Nash equilibrium strategy only matters in treatment info when predicting which strategies subjects will play. In contrast, the coefficients of maxmax and maxmin are highly significant in both treatments. While the three independent variables ("equilibrium", "maxmax" and "maxmin") are correlated, all pairwise correlation coefficients are lower than 0.5. Further details on the relationship of the three independent variables can be found in the appendix.

Result 2. *In both treatments, a strategy is more likely to be played when it cannot lead to the lowest ranked payment pair and when it can lead to the highest ranked payment pair.*

In line with Result 1, the coefficient estimate for the variable "equilibrium" differs significantly among the two treatments and is only useful to predict play in treatment info but not in treatment baseline. In contrast, the highest and lowest ranked payment pair seems to attract our subjects' attention in both treatments. As expected, the according coefficient estimates are higher in treatment baseline than in treatment info. However, the difference is not signficant.[18]

Table 8: Conditional logit regression "played", robust standard errors clustered by subject

Dependent variable: played	Baseline	Info
equilibrium	0.09	0.89****
	(0.21)	(0.22)
maxmax	1.61****	1.17****
	(0.18)	(0.15)
maxmin	1.39****	1.29 ****
	(0.17)	(0.15)
n	1112	1016
Clusters	139	127
Pseudo R^2	0.40	0.41

**** significant at 0.1% level

3.4. Did we manage to elicit subjects' true preferences?

This section first discusses evidence suggesting that subjects did not strategically misrepresent their preferences in stage 1. Then, we show that preference reversals, which might be caused by reciprocal preferences or by the specific game-form, can only lead to a downward bias. Hence, the "true" treatment effect might be even stronger.

Strategic misrepresentation

When preferences are elicited in stage 1 of the experiment, subjects in treatment info are aware that these preferences will be revealed to other subjects. However, they are not informed about the specific games that are played in stage 2. Hence, subjects did not have the information necessary to

[18]The coefficient estimate of an interaction term of maxmin and the treatment dummy (maxmax and the treatment dummy) is not significant.

figure out what kind of misrepresentation might be most advantageous: in some 2×2 games, it could be beneficial to be perceived as having pro-social preferences whereas in other games, the contrary is more likely (e.g., in the chicken game). Moreover, recall that a decision made in stage 2 affects a subject's payment with a probability of only 1/8. All in all, in treatment baseline, subjects had clear and strong incentives to truthfully report their preferences. It was clear to subjects that their reports would not be revealed to anyone but the experimenter. Therefore, it is unlikely that subjects in treatment baseline misrepresented their preferences.

If many subjects in treatment info misrepresented their preferences, we would expect to observe some significant differences between treatment info and baseline. First, we would expect that the preferences reported in stage 1 significantly differ with presumably a higher share of selfish types in the baseline treatment. We checked our data and find that there is no significant difference between the treatments with respect to the most commonly reported preference types. The share of selfish types is even higher in treatment info than in baseline.

Second, and more importantly, we would expect significantly more dominance violations in treatment info than in baseline. We test this hypothesis by using the frequency with which subjects play strictly dominated strategies in stage 2 of the experiment. To identify strategies that are strictly dominated, we use the preferences elicited in stage 1. If these reflect a subject's true preferences, a rational subject should never play such a strictly dominated strategy. In contrast, if subjects strategically misrepresent their preferences in stage 1, a strategy that we classify as strictly dominated may in fact not be dominated according to the subjects' true preferences. Therefore, we can compare the frequency with which subjects play a strictly dominated strategy in the two treatments to test the claim that preferences are truthfully revealed in stage 1 of treatment info. If that claim is true, no difference should be observed. Otherwise, subjects should be more likely to play a strictly dominated strategy in treatment info than in baseline.

Table 6 shows how often subjects play a strictly dominated strategy in the games induced by their preferences reported in stage 1. In order to check the assumption that subjects do not misrepresent their preferences in both treatments, we run a logit regression using the games with strictly dominant strategies (280 in treatment baseline and 295 games in treatment info) as observations. The dependent variable "dominated strategy played" is a dummy variable that assumes a value of 1 if the strictly dominated strategy was played. The only explanatory variable other than the intercept is a treatment dummy ("info") (see Table 9).

Table 9: Logit regression "dominated strategy played", robust standard errors clustered by subject

Dependent variable: dominated strategy played	
Info	0.07
	(0.23)
Constant	-1.20***
	(0.16)
n	575
Clusters	333
Pseudo R^2	0.0002

***significant at 1% level

The coefficient estimate for the treatment dummy is not significantly different from 0. Hence, the null hypothesis cannot be rejected. We also test the same assumption using a two-tailed two-sample Wilcoxon rank-sum test. The dependent variable is then the frequency with which a subject plays a dominated strategy. Each subject who had a strictly dominant strategy in at least one of the four games corresponds to an observation. There are 165 such observations in treatment baseline and 168 in treatment info. We cannot reject the null hypothesis that the frequency with which strictly dominated strategies are played follows the same distribution in the two treatments (p=0.81).

Result 3. *Subjects are equally likely to play a strictly dominated strategy in both treatments.*

Moreover, the fraction of subjects whose reported preferences are consistent with own payoff maximization is even slightly larger in treatment info compared to treatment baseline, though the difference is not significant (p= 0.53 using a Fisher exact test). All other properties of measured preferences that we discussed in section 3.1 are also satisfied equally frequently in both treatments (see tables 4 and 5). We therefore maintain the assumption that subjects truthfully report their preferences in stage 1 of the experiment in both treatments.

Preference reversals

In psychological game theory, Rabin (1993) and Dufwenberg and Kirchsteiger (2004) introduced models of reciprocity in which players reward kind actions and punish unkind ones. Reciprocity could lead to a problem equivalent to the misrepresentation of preferences discussed in this section. For instance, consider Game 1 in stage 2 of treatment info. Suppose an own-payoff maximizer (row) is matched with a total-payoff maximizer (column). The row player might then believe that column will cooperate (play R), even though column expects row to defect (play U). This expected kindness on the part of column might then induce row to also cooperate, thus violating our assumption that only outcomes matter. In other words, subjects' preferences might change once they are shown their opponents' ranking of payment pairs in stage 2 of treatment info. Another potential violation of our assumption might arise if subjects' preferences over payment pairs changed once they are shown the specific game-form.

Since there is a significant treatment effect, such preference reversals would be only a problem if they led to a systematic upward bias (indicating a false-positive result). There is no reason that preference reversals arising from the strategic situation lead to such a bias because subjects face the same game-forms in both treatments. In contrast, reciprocity effects might be relevant here because subjects' preferences are only revealed in treatment info. However, reciprocity effects cannot cause an upward bias: recall that we determined each player's equilibrium strategy based on his reported preferences in stage 1. Furthermore, we only study situations in which each player has a unique equilibrium strategy. If a player changes his preferences, he may want to deviate from this equilibrium strategy. Hence, reciprocal preferences may lead to less (but not more) equilibrium play in treatment info. This suggests a systematic downward bias, which would mean that the "true" treatment effect is even stronger.

We cannot exclude that preference reversals caused some noise, in particular dominance violations, in our experiment. Subjects played in roughly a quarter of the games a strictly dominated strategy (see Table 6). At first glance, this number may appear high. However, it is considerably lower than the number of dominance violations in many standard experiments, which assume that subjects only care about their own monetary payments. For instance, in a meta study by Sally (1995) the proportion of dominance violation in prisoner's dilemma games lies between 30 and 40%.

All in all, we found a significant treatment effect despite many factors that potentially caused noise in the experiment. In our view, this makes our result more robust and indicates that the "true" treatment effect might be underestimated.

4. Conclusion

The assumptions that monetary payoffs in strategic situations represent players' utilities and that their preferences are mutually known are often not satisfied. Our experiment shows that both assumptions play an important role: first, the game that subjects actually play is often very different from the one in which monetary payoffs represent utilities. Second, the treatment effect shows that mutual knowledge of preferences leads to significantly more equilibrium play. Hence, alleged violations of Nash equilibrium can be attributed, to some degree, to the violations of these two assumptions.

The main result of this study suggests that subjects fail to accurately predict other players' preferences and that this significantly affects their behavior. It is plausible that similar difficulties exist in many real-world situations as well because people often have no precise information about how other people evaluate the outcomes of an interaction. Many other models that are used in behavioral game theory (e.g., level-k models) also rely on the assumption of mutual knowledge of preferences. Since in the games we analyzed, subjects often fail to play a Nash equilibrium strategy when preferences are not mutually known, these models also might fail to accurately predict behavior whenever this assumption is not met. Therefore, when deciding what model to apply for analyzing a specific situation, the question whether or not agents can reasonably be expected to know other agents' preferences should play an important role.

Our second result shows that subjects are more likely to play maxmin or maxmax strategies rather than the Nash equilibrium strategy. Hence, many people seem to rely on heuristics rather than on strategic considerations when selecting a strategy.

Acknowledgements. We thank Gary Charness, Jürgen Eichberger, Jörg Oechssler, Christoph Vanberg, the seminar participants at University of Exeter and University of Heidelberg, and the conference participants at ESA

European Meeting 2014, ESA World Meeting 2017, and JE on Ambiguity and Strategic Interactions 2019 for very helpful comments.

References

Attanasi, G., Battigalli, P. and Nagel, R. (2016), 'Disclosure of belief-dependent preferences in a trust game', *Mimeo* .

Aumann, R. and Brandenburger, A. (1995), 'Epistemic conditions for nash equilibrium', *Econometrica* **5**, 1161–1180.

Costa-Gomes, M., Crawford, V. P. and Broseta, B. (2001), 'Cognition and behavior in normal-form games: An experimental study', *Econometrica* **69**, 1193–1235.

Dufwenberg, M. and Kirchsteiger, G. (2004), 'A theory of sequential reciprocity', *Games and Economic Behavior* **47**, 268–298.

Eichberger, J. and Kelsey, D. (2000), 'Non-additive beliefs and strategic equilibria', *Games and Economic Behavior* **30**, 183–215.

Eichberger, J. and Kelsey, D. (2014), 'Optimism and pessimism in games', *International Economic Review* **55**, 483–505.

Ellsberg, D. (1961), 'Risk, ambiguity, and the savage axioms', *Quarterly Journal of Economics* **75**, 643–669.

Fischbacher, U. (2007), 'z-tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics* **10**, 171–178.

Greiner, B. (2015), 'Subject pool recruitment procedures: organizing experiments with orsee', *Journal of the Economic Science Association* **1**, 114–125.

Harsanyi, J. (1967-1968), 'Games with incomplete information played by 'bayesian' players, parts i-iii.', *Management Science* **14**, 159–182, 320–334, 486–502.

Healy, P. (2011), Epistemic foundations for the failure of nash equilibrium. Mimeo.

Hu, A. and Zou, L. (2015), 'Sequential auctions, price trends, and risk preferences', *Journal of Economic Theory* **158**, 319–335.

Lehrer, E. (2012), 'Partially specified probabilities: Decisions and games', *American Economic Journal: Microeconomics* **4**, 70–100.

Lo, K. (1996), 'Equilibrium in beliefs under uncertainty', *Journal of Economic Theory* **71**, 443–484.

McKelvey, R. D. and Palfrey, T. R. (1995), 'Quantal response equilibria for normal form games', *Games and Economic Behavior* **10**, 6–38.

Oechssler, J. and Schipper, B. C. (2003), 'Can you guess the game you are playing?', *Games and Economic Behavior* **43**, 137–152.

Polak, B. (1999), 'Epistemic conditions for nash equilibrium, and common knowledge of rationality', *Econometrica* **67**, 673–676.

Rabin, M. (1993), 'Incorporating fairness into game theory and economics', *American Economic Review* **83**, 1281–1302.

Sally, D. (1995), 'Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992', *Rationality and Society* **7**, 58–92.

Wolff, I. (2014), When best-replies are not in equilibrium: Understanding cooperative behaviour. Mimeo.

Appendix A. (for online publication only)

Appendix A.1. Details of the robustness check tests for the main result

Tables A.10 and A.12 report the results of a two-tailed Fisher exact test of the null hypothesis that the probability that a subject plays the equilibrium strategy is the same in both treatments. These tests were run separately for each of the 4 games. n_base is the number of observations in treatment baseline and n_info the number of observations in treatment info. The tests reported in Table A.10 include all subjects while those reported in Table A.12 include consistent subjects only.

Table A.10: Fisher exact test (two-tailed), all subjects.

Game	n_base	n_info	p-value
1	22	15	0.193
2	11	17	1.000
3	30	32	0.028
4	18	17	0.176
5	17	20	0.082
6	4	3	n.A.
7	22	21	0.364
8	16	14	0.484

Table A.11: Dominance violations per game

Game	baseline	info
1	15/62	16/64
2	3/16	3/19
3	11/38	18/46
4	9/42	9/30
5	9/17	10/25
6	0/5	0/3
7	10/47	11/55
8	8/53	5/53
Total	65/280	72/295

Table A.12: Fisher exact test (two-tailed), consistent subjects only.

Game	n_base	n_info	p-value
1	17	11	0.121
2	7	13	0.374
3	23	25	0.075
4	10	10	0.170
5	17	18	0.075
6	3	2	n.A.
7	22	21	0.364
8	16	11	0.696

Table A.13: Properties of strategies available to consistent subjects, by treatment

equilibrium	maxmax	maxmin	n_baseline	n_info
0	0	0	332	322
0	1	0	236	199
1	1	1	187	189
0	0	1	182	143
0	1	1	61	56
1	1	0	56	44
1	0	1	34	37
1	0	0	24	26

Appendix A.2. Details on the conditional logit regression (table 8)

Table A.13 provides additional information for both regressions that are displayed in table 8: We check for each pure strategy available to consistent subjects whether it is a Nash equlibrium strategy ("equilibrium" = 1), whether it is the maxmax strategy ("maxmax" = 1) and whether it is the maxmin strategy ("maxmin" = 1). "n_baseline" indicates the number of pure strategies in treatment baseline, "n_info" the number of strategies in treatment info.

Appendix A.3. Experiment instructions

Treatment Baseline: Instructions Part 1

1 General Information

Welcome to this experiment and thank you very much for your participation! Please switch off your mobile phone now and do not communicate with each other any more. If you have a question, raise your hand, we will come over to your seat and answer it individually. In this experiment, you can earn a substantial amount of money. The amount you earn depends on your own decisions, the decisions of the other participants and on chance. The amount of money earned will be paid out to all participants individually in cash at the end of the experiment. During the experiment, everyone makes his decisions anonymously on his own. At no point in time will your decisions be linked to your identity.

This experiment consists of two parts, which are identical for all participants: In the first part you are shown eight different payoff-combinations, which you are supposed to evaluate. Each of these combinations consists of two numbers (x, y). The first number x corresponds to the amount of Euro that you receive yourself in this situation. The second number y corresponds to the amount that another participant receives. You are supposed to establish a ranking (a so called "preference relationship") over all these payoff-combinations (x, y). That means, you indicate which of these combinations you like best, which one second-best, and so on. The exact procedure will be explained again step by step later on.

The ranking created in this way, as well your decisions in part two of the experiment, will not be revealed to any other participant. After each participant has created such a ranking over the payoff-combinations, part two of the experiment will begin. Both parts of the experiment are run at the computer. Before they start, you are asked several control questions, which shall help you in your understanding of the experiment. For the second part, you will receive separate instructions. At the end of the experiment, there will be a short questionnaire and then you will be paid in cash.

Your total payoff consists of two payments. In order to determine these payments, one of the decisions made in either part 1 or part 2 of the experiment will be randomly selected. Further details will be provided later on.

2 Evaluation of Payoff-combinations

We will now explain the first part of the experiment, the evaluation of payoff-combinations. You will perform this task immediately afterwards at the computer. You will first be shown the following screen:

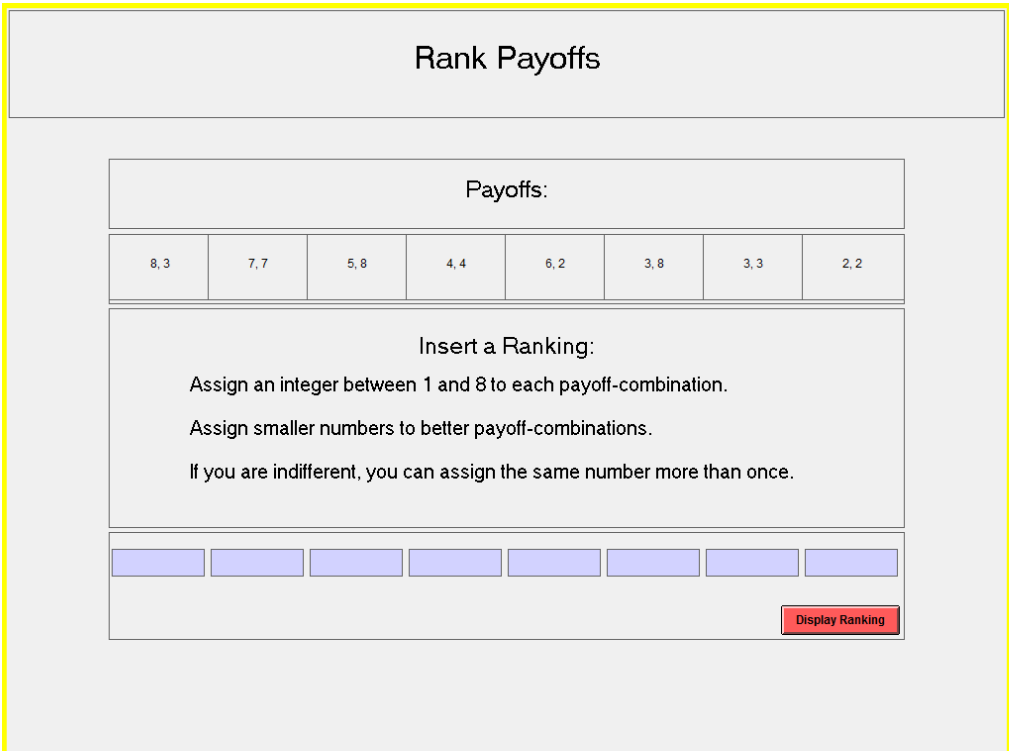

In the row below "Payoffs" you see the eight different payoff-combinations (x, y), which you are supposed to rank (all amounts are in Euro). The payoff combinations are currently ordered randomly. *(Remember: The left value x is the amount you receive yourself and the right value y is given to a randomly selected other participant.)*

You will now assign a number between 1 and 8 to each of these payoff-combinations. The number 1 corresponds to the first rank, which you shall assign to the combination you like best. Analogously the second rank shall be assigned to your second-best combination and so forth until rank 8, which corresponds to your least preferred combination. If you consider two or more combinations as equally good, you are allowed to assign the same rank/number to them.

After you created your ranking, you will see the following screen:

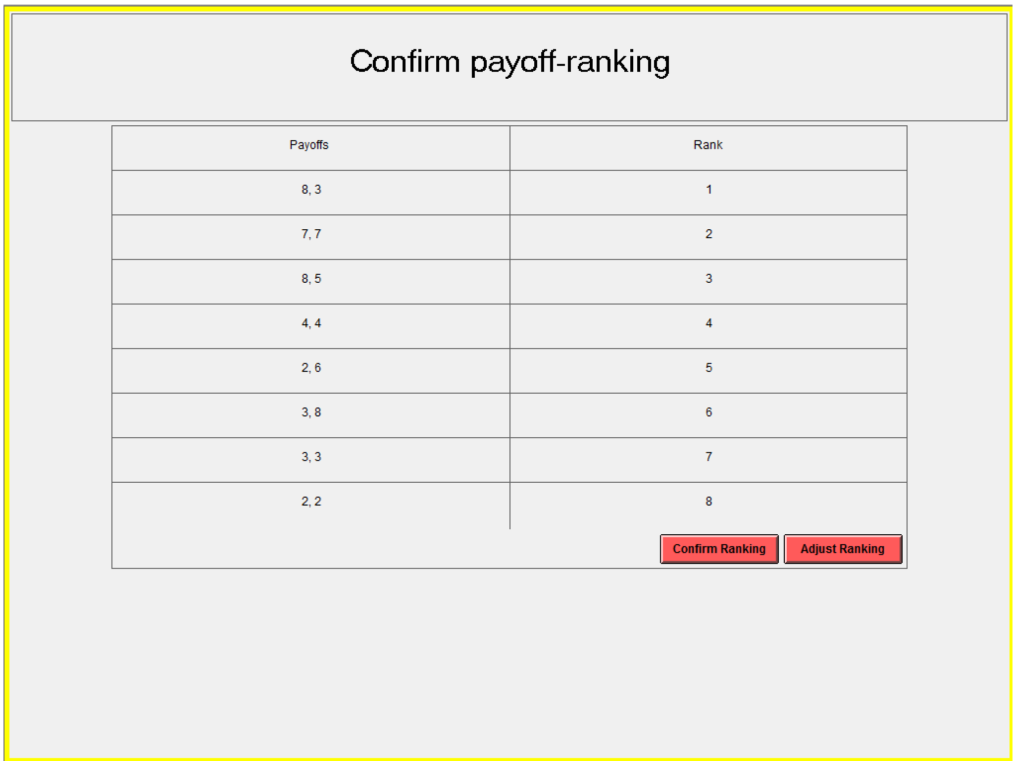

Here you see the payoff-combinations, ordered according to your previously stated preferences. If you like, you can still make modifications. After all participants confirmed their ranking, the second part of the experiment will begin.

3 Calculation of your Final Payoff

The one and only payoff-relevant decision will be randomly selected at the end of the experiment. Your total payoff depends on whether a decision from the first or the second part of the experiment is selected.

With a probability of $\frac{7}{8}$ a decision of part one will be chosen. In this case, two of the eight payoff-combinations will be randomly selected. The payoff combination that you ranked more highly will then be paid out. (If both combinations have the same rank, one of these two will be randomly selected.). You will receive the first amount, the value x. In addition, every participant receives exactly one additional payment y that corresponds to the second amount y of a payment-combination selected for some other participant.*(The assignment is carried out in such a way that the second amount y from your decision is not distributed to a participant you are interacting with during the experiment or from whom you receive the second amount yourself.)*

Payoff, if selected decision is from part one:

Total payoff = Amount x from own decision + Amount y from decision of some other participant

The probability that a decision from part two is chosen for payment is $\frac{1}{8}$. In that case, payments depend on the actions chosen by the participants in part two. The calculation of the final payoff for this case will be explained in the instructions for this part. *(The random draw will be performed by a participant at the end of the experiment. For that purpose he draws a card from a deck containing 32 cards numbered 1 to 32. The numbers 1-28 correspond to all possible combinations of two out of the eight payoff-pairs (x, y)) from the first part. If a number between 29-32 is drawn, a decision from the second part will be paid out.)*

Treatment Baseline: Instructions Part 2

The second part of the experiment is run at the computer as well. This part consists of four strategic decision situations, in the following referred to as "games". In each of these situations, you will be matched with a different participant as game partner, that means you never interact with the same person twice. You and the other player simultaneously select one of two possible actions. The row player always chooses between one of the two actions "up" and "down" and the column player always decides between the actions "left" and "right". *(For the sake of simplicity, the game will be displayed for every participant in such a way, that he always acts in the role as row player and the game partner in the role as column player.)*

In every game, there are four possible outcomes. Which one of these outcomes is selected depends on the action you chose as well as on the action the other player chooses. The four outcomes are are displayed in the form of a payoff matrix. The combination (x, y) in one cell of the matrix corresponds to the amounts of money the two players receive, if the corresponding actions have been chosen. Analogously to the first part, the left value x indicates the amount of money in Euro that you receive and the right value y corresponds to the payoff of the other player. **The combinations (x, y) are chosen in such a way, that they assume the exact same values as those from the first part of the experiment.** Thus in every game there appear four out of the eight payoff pairs evaluated in part one.

If a situation from the second part is chosen for payment, the involved players receive the payoffs that correspond to the outcome of the game. In contrast to the first part, each player only receives one amount of money from the payoff-relevant decision. In addition, each player is given a fixed payment of 5 Euro.

Total payoff = 5 Euro + Payment x obtained in the selected game

In addition to the monetary payments, you are also shown the ranking of the payoff-pairs used in the current game that you submitted in the first part of the experiment.

In the computer program, you will see the following screen:

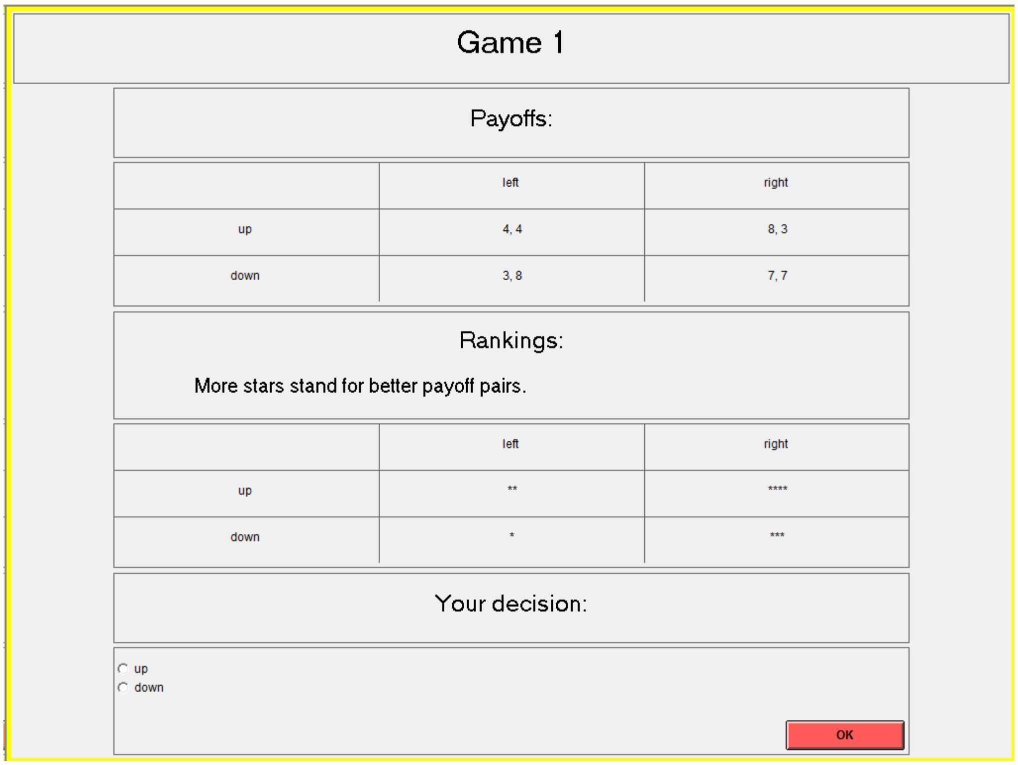

For the sake of clarity, not the exact numbers of the ranking will be shown there, but instead 1-4 stars. A value of four stars (****) means that the corresponding payoff-combination was ranked by you as the best combination (among those appearing in the game). Accordingly, the worst combination is marked by one star (*)

Example:
Let us consider the game shown on the screen "Game 1". If, for example, you decide to play "up" and the other player chooses "right", then you receive a payoff of 8 Euros and your game partner a payoff of 3 Euros. Additionally, you can see in the matrix below, that this is your most preferred outcome.

Are there any questions?

If this is not the case, the second part of the experiment will start shortly...

Treatment Info: Instructions Part 1

1. General Information

Welcome to this experiment and thank you very much for your participation! Please switch off your mobile phone now and do not communicate with each other any more. If you have a question, raise your hand, we will come over to your seat and answer it individually. In this experiment, you can earn a substantial amount of money. The amount you earn depends on your own decisions, the decisions of the other participants and on chance. The amount of money earned will be paid out to all participants individually in cash at the end of the experiment.

This experiment consists of two parts, which are identical for all participants: In the first part you are shown eight different payoff-combinations, which you are supposed to evaluate. Each of these combinations consists of two numbers (x, y). The first number x corresponds to the amount of Euro that you receive yourself in this situation. The second number y corresponds to the amount that another participant receives. You are supposed to establish a ranking (a so called "preference relationship") over all these payoff-combinations (x, y). That means, you indicate which of these combinations you like best, which one second-best, and so on. The exact procedure will be explained again step by step later on.

After each participant has created such a ranking over the payoff-combinations, part two of the experiment will begin. In this part, the information provided in the first part of the experiment will be used. Two participants at a time will be shown each others' ranking of the payoff-pairs provided in part one of the experiment. In both parts of the experiment, you will interact with other participants using a computer. Before we start, you will be asked several control questions, which shall help you in your understanding of the experiment. For the second part, you will receive separate instructions. At the end of the experiment, there will be a short questionnaire and then you will be paid in cash.

Your total payoff consists of two payments. In order to determine these payments, one of the decisions made in either part 1 or part 2 of the experiment will be randomly selected. Further details will be provided later on.

2. Evaluation of Payoff-combinations

We will now explain the first part of the experiment, the evaluation of payoff-combinations. You will perform this task immediately afterwards at the computer. You will first be shown the following screen:

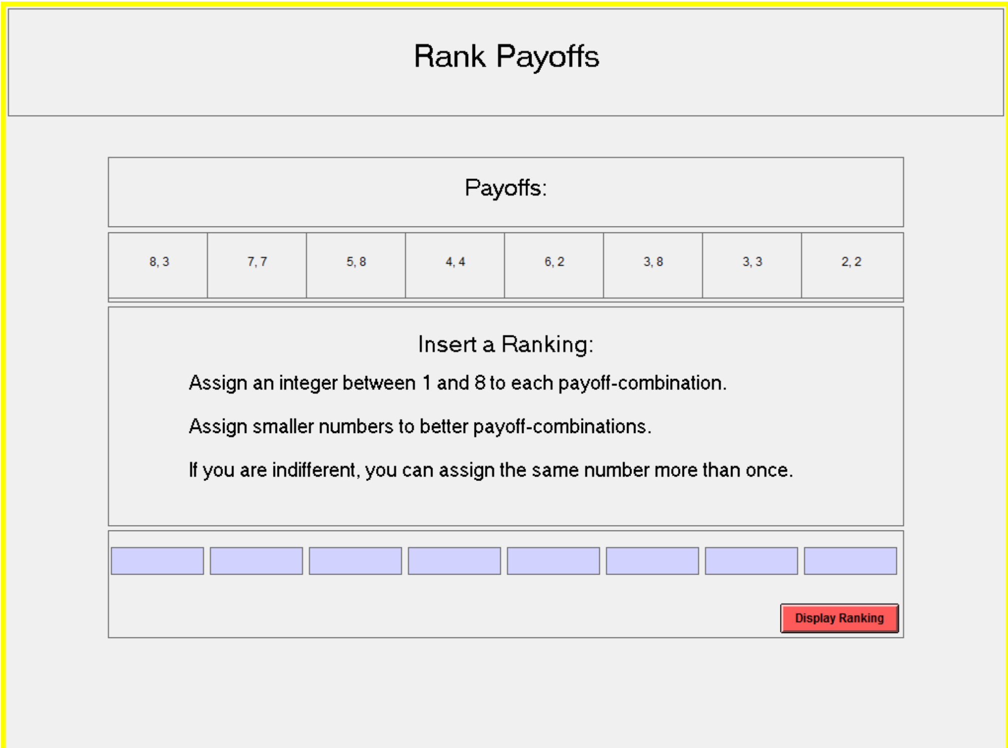

In the row below "Payoffs" you see the eight different payoff-combinations (x, y), which you are supposed to rank (all amounts are in Euro). The payoff combinations are currently ordered randomly. *(Remember: The left value x is the amount you receive yourself and the right value y is given to a randomly selected other participant.)*

You will now assign a number between 1 and 8 to each of these payoff-combinations. The number 1 corresponds to the first rank, which you shall assign to the combination you like best. Analogously the second rank shall be assigned to your second-best combination and so forth until rank 8, which corresponds to your least preferred combination. If you consider two or more combinations as equally good, you are allowed to assign the same rank/number to them.

After you created your ranking, you will see the following screen:

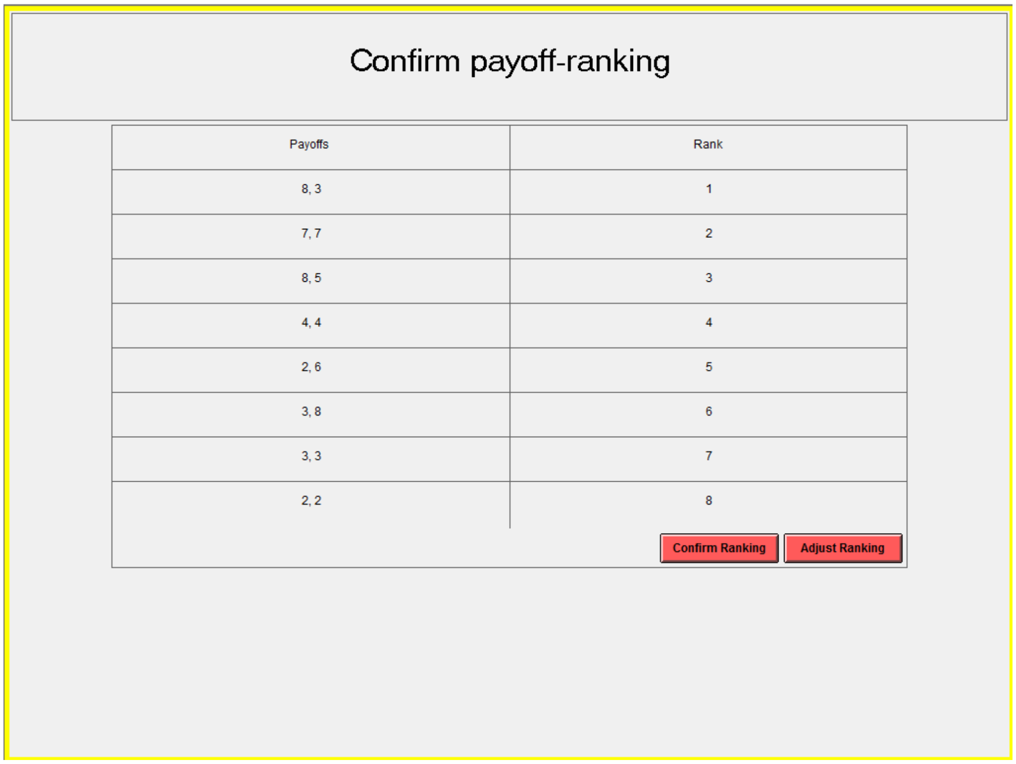

Here you see the payoff-combinations, ordered according to your previously stated preferences. If you like, you can still make modifications. After all participants confirmed their ranking, the second part of the experiment will begin.

3. Calculation of your Final Payoff

The one and only payoff-relevant decision will be randomly selected at the end of the experiment. Your total payoff depends on whether a decision from the first or the second part of the experiment is selected.

With a probability of $\frac{7}{8}$ a decision of part one will be chosen. In this case, two of the eight payoff-combinations will be randomly selected. The payoff combination that you ranked more highly will then be paid out. (If both combinations have the same rank, one of these two will be randomly selected.). You will receive the first amount, the value x. In addition, every participant receives exactly one additional payment y that corresponds to the second amount y of a payment-combination selected for some other participant.*(The assignment is carried out in such a way that the second amount y from your decision is not distributed to a participant you are interacting with during the experiment or from whom you receive the second amount yourself.)*

Payoff, if selected decision is from part one:

Total payoff = Amount x from own decision + Amount y from decision of some other participant

The probability that a decision from part two is chosen for payment is $\frac{1}{8}$. In that case, payments depend on the actions chosen by the participants in part two. The calculation of the final payoff for this case will be explained in the instructions for this part. *(The random draw will be performed by a participant at the end of the experiment. For that purpose he draws a card from a deck containing 32 cards numbered 1 to 32. The numbers 1-28 correspond to all possible combinations of two out of the eight payoff-pairs (x, y)) from the first part. If a number between 29-32 is drawn, a decision from the second part will be paid out.)*

Treatment Info: Instructions Part 2

The second part of the experiment is run at the computer as well. This part consists of four strategic decision situations, in the following referred to as “games”. In each of these situations, you will be matched with a different participant as game partner, that means you never interact with the same person twice. You and the other player simultaneously select one of two possible actions. The row player always chooses between one of the two actions “up” and “down” and the column player always decides between the actions “left” and “right”. *(For the sake of simplicity, the game will be displayed for every participant in such a way, that he always acts in the role as row player and the game partner in the role as column player.)*

In every game, there are four possible outcomes. Which one of these outcomes is selected depends on the action you chose as well as on the action the other player chooses. The four outcomes are are displayed in the form of a payoff matrix. The combination (x, y) in one cell of the matrix corresponds to the amounts of money the two players receive, if the corresponding actions have been chosen. Analogously to the first part, the left value x indicates the amount of money in Euro that you receive and the right value y corresponds to the payoff of the other player. **The combinations (x, y) are chosen in such a way, that they assume the exact same values as those from the first part of the experiment.** Thus in every game there appear four out of the eight payoff pairs evaluated in part one.

If a situation from the second part is chosen for payment, the involved players receive the payoffs that correspond to the outcome of the game. In contrast to the first part, each player only receives one amount of money from the payoff-relevant decision. In addition, each player is given a fixed payment of 5 Euro.

Total payoff = 5 Euro + Payment x obtained in the selected game

As announced before, you will now receive information about each others’ preferences. This means that in addition to the payoff matrix you are shown another matrix below, in which you can see how you and the other player ranked the payoff combinations used in the current game in the first part of the experiment. *Note: you interact with a different partner in every game and therefore the ranking of your opponent may change from one game to another.*

In the computer program, you will see the following screen:

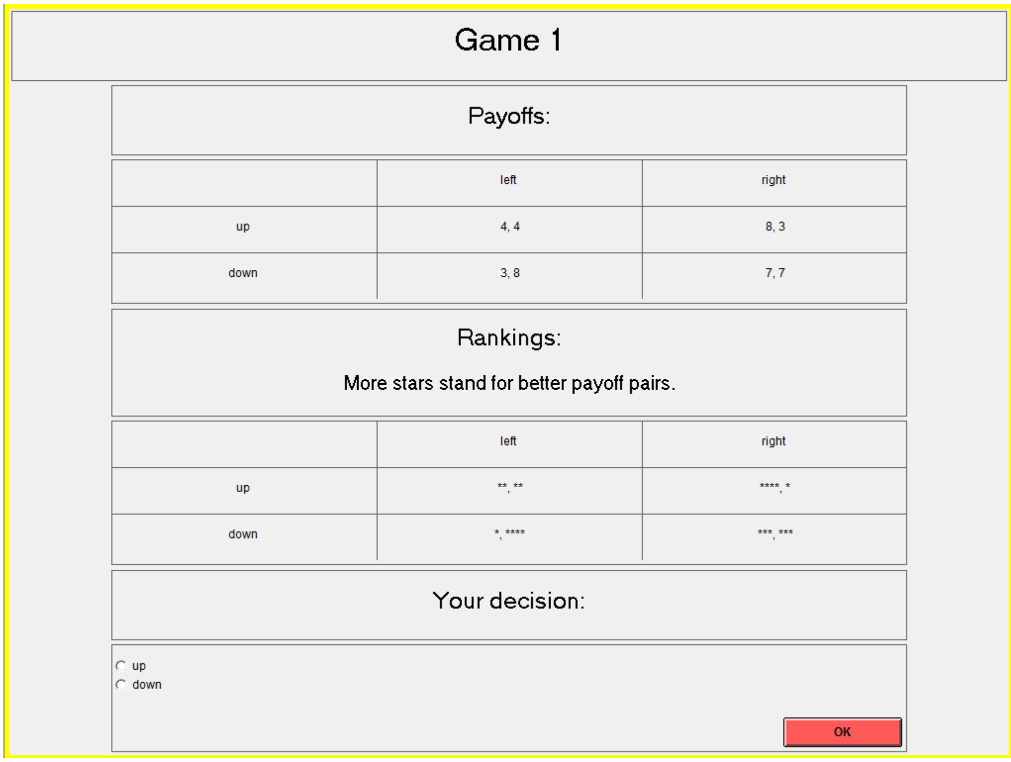

For the sake of clarity, not the exact numbers of the ranking will be shown there, but instead 1-4 stars. A value of four stars (****) means that the corresponding payoff-combination was ranked by you as the best combination (among those appearing in the game). Accordingly, the worst combination is marked by one star (*)

Example:
Let us consider the game shown on the screen "Game 1". If, for example, you decide to play "up" and the other player chooses "right", then you receive a payoff of 8 Euros and your game partner a payoff of 3 Euros. Additionally, you can see in the matrix below, that this is your most preferred outcome, but the least preferred outcome of the other player.

Are there any questions?

If this is not the case, the second part of the experiment will start shortly...